

AD-A183 832

TOWARD A THEORY OF CURRICULUM FOR USE IN DESIGNING
INTELLIGENT INSTRUCTION (U) PITTSBURGH UNIV PA LEARNING
RESEARCH AND DEVELOPMENT CENTER A M LESGOLD

1/1


UNCLASSIFIED

83 AUG 87 UPITT/LRDC/ONR/LSP-2


F/G 12/9

NL

END
8/87
DTTC



AD-A183 832



University of Pittsburgh

12

DTIC FILE COPY

Toward A Theory of Curriculum For Use in Designing Intelligent Instructional Systems

Alan Lesgold

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTED
AUG 31 1987
S D
CD

August 3, 1987

Technical Report No. LSP-2

This research was sponsored by the Psychological Sciences Division, Office of Naval Research,
under Grant No. N00014-83-K-0655, NR 667-524.

Reproduction in whole or part is permitted for any purpose of the United States Government.

Approved for public release; distribution unlimited.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

A183 832

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) UPITT/LRDC/ONR/LSP-2		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Learning Research & Development Ctr, U. of Pittsburgh	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Personnel & Training Research Programs Office of Naval Research (code 1142PT)	
6c. ADDRESS (City, State, and ZIP Code) Pittsburgh, PA 15260		7b. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, VA 22217-5000	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0655/P00004	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO. 61153N	PROJECT NO. RR04206
		TASK NO. RR04206-00	WORK UNIT ACCESSION NO. NR442c524
11. TITLE (Include Security Classification)	Toward a Theory of Curriculum for Use in Designing Intelligent Instructional Systems (UNCLASSIFIED)		
12. PERSONAL AUTHOR(S)	Lesgold, Alan M.		
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) April 1987	15. PAGE COUNT
16. SUPPLEMENTARY NOTATION A version of this report will appear as a chapter in Mandl, H. and Lesgold, A., (Eds), "Learning Issues for Intelligent Tutoring Systems," Springer Verlag, N.Y.			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Artificial intelligence; Cognitive psychology; Cognitive science; Instruction, computer-based; Instruction, intelligent computer-based; Training, technical; Tutors; Tutors, intelligent.	
FIELD 05	GROUP 09	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) In general, intelligent computer-based tutors have an explicit representation of the knowledge to be taught but lack an explicit theory-based representation of the curriculum, the goal structure for teaching that knowledge. Adding goal structure representations would permit a richer approach to such problems as remediation. This paper addresses such questions as: When prerequisite subgoals have been satisfied but a current goal fails, how do we remediate? When subgoals are satisfied, do we know the trainee will know when to apply the subskills? We propose an architecture for intelligent tutors that explicitly represents target (expert) knowledge; curriculum or goal knowledge; and knowledge of the individual trainee's aptitude.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Susan M. Chipman		22b. TELEPHONE (Include Area Code) (702) 696-4318	22c. OFFICE SYMBOL ONR 1142PT

**TOWARD A THEORY OF CURRICULUM
FOR USE IN DESIGNING
INTELLIGENT INSTRUCTIONAL SYSTEMS**

Alan Lesgold

**Learning Research and Development Center
University of Pittsburgh
Technical Report No. LSP-2**

This research was sponsored by the Psychological Sciences Division, Office of Naval Research, under Grant No. N00014-83-K-0655, NR 667-524. The report refers to collaborative research involving Jeffrey Bonar, Marilyn Bunzo, Cindy Cosic, Robert Cunningham, Marty Kent, Susanne Lajoie, Debra Logan, Mary Ann Quayle, Peter Reimann, Paul Resnick, Valerie Shute, William Weil, Leslie Wheeler, and others, under the sponsorship of the Office of Naval Research, the Air Force Human Resources Laboratory, and the National Institute of Education. The opinions expressed are solely the author's; none of the funding agencies or collaborators named necessarily endorses or agrees with the views expressed.

Reproduction in whole or part is permitted for any purpose of the United States Government.

Approved for public release; distribution unlimited.



Accesion For	
NTIS	CRA&I <input checked="" type="checkbox"/>
DTIC	TAB <input type="checkbox"/>
Unannounced <input type="checkbox"/>	
Justification	
By	
Distribution /	
Availability Codes	
Distr	Avail and/or Special
A-1	

Implicit in the approaches being taken by current efforts to create intelligent computer-based instruction is the notion that curriculum is almost an epiphenomenon of knowledge-driven instruction. Early computer-based instruction had little control structure other than an absolutely rigid curriculum and was insensitive to the subtleties of different students' partial knowledge. As a result, there was a reaction in the direction of representing the students' knowledge as a subset of the target or goal knowledge to be taught and simply deciding de novo after each piece of instruction what piece of missing knowledge to teach the student. I am convinced that goal knowledge is as important to intelligent machine activity as it is to human activity, and that it also must be well understood and explicitly represented in an instructional system if that system is to be successful in fostering learning.¹ This report presents an architecture for representing curriculum or goal knowledge in intelligent tutors and is thus a first step toward a theory of curriculum that can inform the design of such systems. To illustrate one way in which such a theory can sharpen our ideas about learning and instruction, the later part of the report focuses on the concept of prerequisite that is the basis for existing computer-assisted instruction and shows how that concept has been inadequate in the past. A new approach, in which the prerequisite relationship is always dependent on the instructional subgoal (curriculum) context, is introduced.

1.0 Current Practice

Programs that preceded the entrance of artificial intelligence into instruction prespecify the content of lessons. In some cases, the order of the specific lessons to which a student is exposed is computed as instruction proceeds. However, the content of a lesson, in terms of knowledge it is trying to teach, tends to be fixed. In currently used programs, lesson assignment is viewed as a more-or-less knowledge-free subgoaling problem. We have a list of things to be taught, and we teach each in turn. If instruction is unsuccessful, we try again. The assignment of lessons can occur in several different

ways. The *frame* approach organizes instruction into very microscopic units, characteristically one screen "frame" each in size. Each frame contains some instruction, some means for testing the student's performance, and decision rules for deciding, on the basis of that performance, which frame the student should see next. The *pretest-posttest* approach is organized at a higher level, the level of the lesson. Prior to each lesson, a pretest is given. If the pretest is passed, the lesson is skipped; otherwise it is presented. After each lesson a posttest is given. If the student passes the posttest, he goes on to take the pretest for the next lesson.

There is a certain amount of inefficiency in the pretest-posttest approach. In many cases, the student spends most of his time taking tests, and many testing methods are not very effective instruction. More important, there usually is not much difference between the lesson taken initially by a student and the one he receives if he fails the posttest and is recycled. A number of efforts have been made in specific programs to assure that novel material is presented, or that the repeated lesson is taught more slowly, with more examples and more practice. However, in terms of content, it is the same old lesson, being repeated again. As we shall see, this may be a fundamentally incorrect approach to teaching.

There are many rationales for the pretest-posttest approach. Of these, perhaps the strongest is the learning hierarchy theory of Robert Gagné (1962). Gagné gave very clear directions for deciding on the content and sequencing of instruction, and these directions continue to have strong influence on the design of instruction and training today. The basic approach is to start with the capability that is the goal of the training being designed.² One is then to ask the question

What kind of capability would an individual have to possess if he were able to perform this task successfully, were we to give him only instruction? (Gagné, 1962, p. 356)

That is, what does the trainee have to know so that simple verbal advice is sufficient to get him to apply that prior knowledge to the task at hand? This methodology can, of course, be applied recursively to further decompose the target capability into smaller and smaller prerequisite capabilities. It is, after all, nothing but the generic problem solving method of progressive refinement, splitting a complex goal into several pieces, splitting the pieces into still smaller pieces, etc. This approach works as long as we pick as units pieces that are small and coherent enough to cover in single lessons and as long as there are no dangerous interactions between achieving one lesson subgoal and achieving another.

These two criteria of subgoal internal coherence and linearity are extremely important to the success of the method of progressive refinement. By internal coherence, we mean that the subgoal can be achieved sensibly by itself, without duplicating effort across subgoals unnecessarily. For example, if we decided to mow the lawn by dividing it into a checkerboard of square regions, first mowing every other region, and then mowing the ones in between, our refinement of the lawn task into doing first the red squares of the board and then the black ones wastes effort. We must move over the entire lawn twice instead of once. Further, lifting the mower to avoid cutting certain sections which are being postponed for later is also unnecessary work.

Note that internal coherence is a function of the specifics of subgoal contents and not a general principle alone. For example, while the checkerboard approach to mowing the lawn is inefficient, making two passes over the lawn, one to mow and one to remove weeds, may be quite sensible. This is because the micro-acts of pulling one weed and mowing a small patch are incompatible, producing extra thinking and physical work if they must be continually alternated, while the micro-acts of mowing successive small regions are quite compatible, saving work over doing one patch, then mowing it, doing the next, etc. It might be argued that the current approach of teaching subtraction of

two-digit numbers and then waiting as long as a year before taking up subtraction of three-digit numbers also violates the internal coherence rule.

The linearity criterion is also important. For example, if we want to build a house, we can split the task into laying a foundation, putting up walls, and putting on a roof. Only one order of those three steps make sense. The other orders fail. Some orderings are impossible: one cannot put on a roof if there are no walls to support it yet. Other orderings involve early steps that interfere with later ones, e.g., putting up the walls before the foundation. Linearity problems can also arise in learning. For example, if one requires students to type their essays beginning in fifth grade, and typing is not taught until seventh grade, then there will be some conflict. Students will acquire, on their own, patterns of typing that conflict with the behaviors used in efficient typing methods; there could be negative transfer.

Gagné (1971) undertook to prescribe some principles of learning that would help instructional designers achieve internal coherence and linearity in their development of hierarchical goal structures (which he called learning hierarchies) for training courses. He developed a variety of specific learning forms that could be used to constrain the parceling of pieces of instruction into separate lessons or curriculum subgoals. Implicit in his work is the principle that, in learning, the whole is more than the sum of its parts. That is, the lowest-level subgoals in a goal hierarchy do not, as a group, contain all the knowledge implied by the highest-level goal. The instructions that are given as subgoals are assembled into larger units of capability that result in new learning. This new learning is not part of any subgoal's knowledge. Rather, it is emergent when multiple subgoals are combined, just as a theorem in geometry is not present in the premises from which it is derived but is, rather, new knowledge.

Unfortunately, the belief that the whole is simply the sum of its parts has too often guided the task analyses that have generated training curricula. Rational task analysts

have tended to look at an instructional subgoal, intuitively decide whether it can be taught in a single lesson, and, if not, split it into a few subsubgoals each to be treated by a separate lesson. Gagné's criteria, that we should be able to provide the "glue" needed to tie together the pieces of knowledge from prerequisite lessons using only simple instructions, and that these instructions should be based upon a theory of learning, are generally not given adequate consideration. For example, designers of training for technicians make intuitive decisions that a good technician must know some of the theory of operation for the devices he maintains and that he must know some specific rules for fixing those devices. A current technical curriculum therefore may cover primarily decontextualized theory and device-specific operating algorithms. No effort appears to have gone into determining the cognitive glue that allows concepts learned in one context to be applied in another, such as finding a fault in a nonworking device, which may require problem solving heuristics rather than specific algorithms.

Intelligent computer-assisted instruction attempts to represent all the knowledge that constitutes the expertise that is to be taught. Interestingly, though, it generally does not possess an explicit curriculum based upon a theory of learning and instruction, either. Where conventional instruction has an explicit curriculum but fails to have an explicit and complete representation of the knowledge that is to be taught, intelligent instructional systems have tended to represent the target knowledge explicitly but not to represent explicitly that body of knowledge that specifies the goal structure for instruction, the curriculum. For example, the WEST tutor (Burton & Brown, 1982) contains a method for determining how close to optimal a player's performance is and a set of issues to be considered. These issues constitute part of a curriculum knowledge structure but fail to have any relational structure tying them to each other or to a representation of target knowledge. In other cases, such as the geometry and lisp tutors being developed by John Anderson (Anderson, Boyle, Farrell, & Reiser, 1984), there are

problem sequences which are preset, but again there is no explicit representation of curriculum knowledge.

To summarize, traditional computer-based instruction, whether organized into frames or into larger lesson units, tends to have an explicit representation of curricular structure, though often a shallow one, and, at best, only an implicit representation of the knowledge being taught, while intelligent instructional systems developed to date have explicit representation of the target knowledge but at best only implicit representations of the curriculum knowledge, the scope and sequencing of lessons.

2.0 Weaknesses in Current Approach

We can now restate and elaborate two problems inherent in current approaches to instructional design. First, there is no clear method for differentiating how to present material to remediate a problem discovered after a lesson has been taught from how it should be presented when taught initially. Second, the knowledge that represents the "glue" connecting the contents of related lessons is not clearly specified, nor is it assigned to be part of the content of any specific lesson. I consider each of these problems in turn.

2.1 *Redo is the Only Strategy It Supports*

In current training systems, the curriculum is at least implicitly a goal structure. One proceeds to teach the prerequisites (or subgoals) for a given lesson before teaching that lesson itself. The student should always be able to infer the missing knowledge that integrates those pieces of prerequisite knowledge into a broader skill. However, this doesn't always work. Lesgold's Two Fundamental Laws of Instruction (shown in Table 1) often apply, resulting in incorrect decisions about whether or not a trainee has mastered a given lesson. A lesson can appear to be mastered although the knowledge that has been acquired is too specific and cannot transfer from the context of the prerequisite training

to the context in which it must be used (missing "glue"), resulting in the circumstance described by the First Law. On the other hand, a lesson can appear to be unlearned because the context of testing does not match the context of instruction, even if the learned knowledge is adequate to the contexts for which the lesson is prerequisite. The result is the circumstance described by the Second Law. For example, there are mechanics whose formal knowledge of electrical principles is not sufficiently developed to pass tests but who know more about what goes on in a car's electrical system than many people who have studied physics.

Table 1

Lesgold's Two Fundamental Laws of Instruction

First Law: Not everyone who passes a test on a topic knows what appears to have been tested.

Second Law: Not everyone who fails a test on a topic lacks the knowledge that appears to have been tested.

Because we are never able to establish with certainty that a lesson has been learned, and because excessively high criteria for posttesting can be very wasteful of instructional time (since the student could fail but still have adequate knowledge), it is inevitable that occasions will arise in which it was assumed that a student has mastered a prerequisite when in fact he has not. It is important to note that this will occur whether the criteria for passing out of a lesson are prespecified or determined through some inferential process as the lesson is being taught, whether they are superficial (a cutoff score on a test) or deep (based upon a detailed student model fitted to all of the student's performance in the relevant recent past).

When the student does have to be given remediation because the assumptions of prerequisite knowledge have proven wrong, current systems generally replay the same instruction that did not work the first time. Sometimes different problems are assigned as examples or for practice, but they are generally of the same type as were used before. Sometimes the lesson proceeds more slowly, taking smaller steps and providing more practice at each step. In a few of the most recent intelligent computer assisted instructional systems (e.g., Burton & Brown, 1982; Bonar, 1985), it is possible to estimate which specific pieces of knowledge targeted by the lesson are most likely not to have been learned and to concentrate on those. However, in every case, the goals of the lesson can and should be adapted to specific needs that arise only during remediation.

It should be noted that a number of theorists have proposed the view that learning in a domain is a process of successively replacing primitive conceptions, or personal theories, of the domain with more advanced constructions of it. This view subsumes the important insight that the knowledge structure of the student, rather than simply being incomplete, may actually be wrong, that he may hold a misconception, a different and conflicting theory of the task domain from the one the trainer would like him to have (for example, see Carey, in press; Glaser, 1984; Shaughnessy, 1977; Young & O'Shea, 1981). The approach I am taking in this report is complementary to this view, concentrating on a somewhat more microscopic level of analysis.

2.2 No One Is Responsible for the Likeliest Failures of Instruction

A second problem with current approaches to remediation is that they fundamentally ignore the nonlinearities between lessons, the "glue" that holds lessons together. This is a problem particularly when responsibility for different parts of a training regimen is divided among multiple instructors. There is always a tendency for the content of a lesson to be abstracted to its core, both in teaching it and in deciding how to test it. Consequently, the amount of between-lesson "glue" for which none of the

course instructors takes responsibility can be substantial. Even worse, such division of responsibility is most common in technical training courses, where the trainees are often less likely to be facile at inferring the missing content. However, it occurs between grades in school and between courses in college curricula as well.

The loss of the fringe content between lessons, the "glue," produces a variety of pathologies that we see every day. The teacher in third grade feels that the only reason students do poorly in her class is that the second grade teacher failed to teach what was required. The second grade teacher points to high test scores and disclaims responsibility. The trainer providing on-the-job practice to technicians claims that they do poorly only because they were not taught fundamental principles of electricity, while the instructor for the course on fundamentals of electricity has lots of test data to show that they learned everything in the curriculum. Of course, what they were likely to forget after testing is exactly the fringes of the knowledge they were taught, the relations that "glue" it to the content of other lessons.

Overall, both of the above-cited problems seem to arise because there is a lack of distinction between the content of training and the curriculum or goal structure for training. Merely checking off subgoals as they are taught fails to take account of the tendency of declarative knowledge to suffer high forgetting at its fringes, to shrink to a coherent, highly-interconnected core. Simply reteaching prerequisite lessons when problems arise later fails to take account of what has already been learned and what has just been shown specifically to be weak.

In the next section, I introduce a knowledge architecture for intelligent tutoring systems that has driven considerable current work on intelligent instructional systems that Jeffrey Bonar, Robert Glaser, and I (cf. Bonar, 1985; Glaser, Lesgold, & Lajoie, in press) have been conducting. The architecture's components are still being shaped by

various projects each of us is conducting, but the basic ideas are well enough evolved to help supply what is missing from current instructional systems.

3.0 The Structure of Knowledge in an Intelligent Tutor.

As discussed above, a fundamental problem with existing architectures for instructional and training systems is that they fail to explicitly represent either the knowledge that they are designed to teach or the curriculum (goal structure) for teaching it. In setting out to develop an architecture that represents both content and curriculum, one quickly discovers that there are other issues that must be considered. First, there is other knowledge that a good teacher looks for in the student, knowledge we might call aptitude or metacognitive skill. That is, some students are more able to learn with facility from particular forms of instruction than others, and awareness of a student's learning capabilities can well shape the instructional approach a good teacher takes. Second, there are more domain-specific capabilities that represent both the broad outcomes desired from the training and some specific capabilities for learning. For example, in a course on troubleshooting, general knowledge of electrical principles is both a useful prerequisite and something that should be enhanced by practice in finding faults in circuits.

These needs suggest that the knowledge in an intelligent tutor must be of three different types: (a) curriculum knowledge, a subgoal lattice of lessons connected by the prerequisite relation; (b) a representation of the knowledge to be taught, from which explanations and student models can be generated; and (c) a representation of the more enduring characteristics (metaissues) to which the instruction should be sensitive. Figure 1 shows the architecture symbolically.




Figure 1
The Three Layers of an Intelligent Tutor

3.1 The Knowledge Layer

The knowledge layer, highlighted in Figure 2, should contain a representation of the knowledge the system is trying to teach. One way to think about that knowledge is that

it is a model of expert capability in the domain. Such knowledge includes both procedures and concepts (i.e., both procedural and declarative knowledge). Some constraints on its structure can be inferred from several things we know about human expertise. For example, we know that experts, in contrast to novices, tend to represent problems according to the underlying situations they involve, according to their deep structure, whereas novices tend to have more superficial representations (Chi, Feltovich, & Glaser, 1981; Larkin, McDermott, Simon, & Simon, 1980). We also know that, in contrast to intermediate-level performers, experts tend to know exactly what to do in a given situation rather than being dependent on inference from first principles (Chase & Simon, 1973; de Groot, 1965). That is, they are able to represent the situation more completely and richly and then to invoke the precisely appropriate method for dealing with it. In order to do this, experts' knowledge must be richly interconnected and, to some extent, redundant.

Before discussing how to deal with this problem, let me pause and give an example or two to illustrate the importance of the bridging connections between coherent bodies of knowledge. Consider the field of medicine. It is driven by several sciences: physiology, biochemistry, pharmacology, and even physics and chemistry. The relative roles of different portions of its scientific backing will differ for different disease problems. For this reason, medicine is organized into specialties, each of which is internally very coherent. The ties between these specialties are much more complex and, relatively speaking, ad hoc.




Figure 2
The Knowledge Layer

A friend who is an attorney is handling a malpractice case involving a man who arrived at a hospital with severe back pain and several other symptoms. A very inexperienced intern was in the emergency room. He sent the patient to the orthopedics department. The specialists there found no skeletal problems and, after a lengthy examination, sent him back to the emergency room again. It then occurred to the intern that the man might have a tumor producing the pain. So, he was sent to oncology. While waiting for his turn there, he collapsed and died of a renal artery aneurysm. In essence, the diagnosis process, decentralized according to the primary joints in the body of

medical knowledge, failed to make adequate use of the linking knowledge that might have led from the symptoms, which seemed to involve the back or the organs near the painful spot, to knowledge of various sorts of vascular problems. What was missing from the process was exactly the ad hoc interconnections between the conventional units of knowledge. Part of the specialty of emergency medicine is knowledge of the limits of the diagnostic processes which different specialties use, i.e., knowledge of what happens at the "fringes" between specialized diagnostic approaches.

One can find other less-dramatic examples throughout the curriculum. For example, in teaching elementary arithmetic, we teach children about place value, concentrating on the ones, tens, and hundreds places in numbers. We also teach them algorithms for addition and subtraction of multicolumn numbers. Not all of the ties between these two related pieces of knowledge are explicitly taught, and not all seem to be universally learned (as demonstrated by the BUGGY line of research; Brown & Burton, 1978; Brown & Van Lehn, 1980; Van Lehn, 1983). Some of what is not universally learned also has the character of being in the gap between the two pieces of instruction.

How do we deal with such gaps? One approach is a small extension of the original Gagné learning hierarchy ideas. The lowest levels of lessons in such a hierarchy correspond to the regions of the knowledge layer into which the total body of expert knowledge has been split. Higher levels of lessons are more than just the sum of what was taught in the lower-level lessons. They have the specific task of assuring that the conceptual glue between the lower-level pieces is acquired. It is in this sense that teaching the whole of a body of material is more than just teaching its parts; the goal for the whole includes not only the parts but also a specific focus on the ties between those parts.




Figure 3
Goal Hierarchy for Basic Resistor Network Laws

3.2 The Curriculum Goal Lattice Layer

The curriculum goal lattice layer is the central layer of the proposed architecture. As the goal structure for the instructional system, it is, more or less, in control of the system. Ordinarily, goal structures are trees, representing the progressive decomposition of each layer of subgoals into still smaller subsubgoals. An example tree is shown in Figure 3. It shows the decomposition of a basic course in resistor network concepts into two main goals, knowledge of Ohm's Law and Knowledge of Kirchhoff's Law. Those goals are then broken into subgoals, which are then decomposed further. The diagram becomes rather complex visually, but its underlying structure is still straightforward: each subgoal is either a lesson that can be taught completely as a unit or it is further decomposable into subsubgoals.

This kind of goal structure is exactly the sort of concept that Gagné was introducing in his discussions (see above) of learning hierarchies. Further, he felt that psychological laws of learning would determine when subgoals had to be further subdivided and when they could be taught as single lessons. His analyses were in terms of the verbal association theory prominent at the time. More recently, Van Lehn has advanced at least one different sort of criterion for deciding on such a subdivision, namely that a single lesson should not require the student to learn a rule with disjunctive conditions (Van Lehn, 1983). Clearly, if the structure of curriculum is simply a subgoal tree, we are well on the way to understanding how to develop such a tree and how detailed its arborization must be.

However, in the first efforts, by Jeffrey Bonar, his students Cynthia Cosic and Leslie Wheeler, and I, to employ a curriculum goal hierarchy in an intelligent tutor, things were not as simple as we had hoped. For example, while one valid way to think of the resistor networks course is in terms of scientific laws presented, which leads to the decomposition shown in Figure 3, there are other equally valid ways. For example, one

might start with the basic measurable properties of such networks: current, electromotive force (voltage), and resistance. This leads to a goal lattice such as that shown in Figure 4. What is noteworthy is that the lowest level units in the tree, the simple lessons, are the same as in Figure 3, but the organization into higher-order goals is entirely different, and the apparent purpose of the course may be different.




Figure 4
Goal Hierarchy for Basic Resistor Network Measures

So far, in our own work, we have found four different viewpoints on the instruction that we want to present, each of which gives rise to a hierarchy that projects onto the same simple lessons. Figure 5 shows these viewpoints. We can partition our lessons into those that deal with series circuits and those that deal with parallel circuits, quite reasonable given that students often have different conceptual problems with parallel circuits. Or, we can partition our lessons according to the type of problem we present to the student: qualitative problems, quantitative problems, and problems that involve making a relative judgment about the one circuit relative to another. There are, of course, also the two viewpoints discussed above, laws and measurable properties.




Figure 5
Viewpoints on Resistor Network Instruction

This leads us to a new view of the structure of curriculum knowledge in the knowledge base of an intelligent tutor. The curriculum knowledge has the structure of a *goal lattice*. There are a number of viewpoints on the goals of the instruction. With respect to each viewpoint, one can identify a subset of the curriculum lattice that is a true subgoal tree structure. So, from any specific point of view, there are clear pathways that determine the sequencing of instruction, though of course there are alternate approaches to such sequencing. For example, one can proceed depth-first. In

the case of Figure 4's viewpoint, this would mean perhaps doing all the lessons relating to current, then all those relating to voltage, and finally all those relating to resistance. Or, one might proceed breadth-first, going through all the lowest-level lessons, then the next level, and so on. There may be individual differences in aptitude or preference for these two approaches.

Of course, when all of the viewpoints are considered at once, there is much more complexity to the task of deciding what the appropriate sequencing for the lessons of the curriculum should be. To some extent, the decision can be made on the basis of rational task analysis, but our experience has been that empirical work driven by cognitive theory is often necessary (Lesgold et al., 1986). Some of the lessons will tend to be difficult and others easy. By taking the various viewpoints, it should be possible to organize knowledge about lesson difficulty sufficiently to use it in deciding on appropriate orderings through the curriculum. Another approach may be to tell the student which lessons he is "eligible" to take next, based on prerequisites completed, and let him decide for himself. As we shall see in section 4 of this report, there are even more sophisticated possibilities to consider in deciding how to handle sequencing.

To summarize, the goal lattice layer is a lattice structure in which are embedded a number of goal hierarchies, each corresponding to a fundamental viewpoint on the task of teaching the course content. Figure 6 shows the goal lattice for the resistor networks course we are implementing. This multiple viewpoints approach, incidentally, has implications for what constitutes an appropriate course, in terms of the completeness, coherence, and consistency of its curriculum lattice. Presumably, the resistor networks course as shown in Figure 6 is a reasonably sensible selection of content for a course. The course is coherent, in that each simple lesson is relevant to all of the viewpoints we have taken. It is locally complete, in that each viewpoint seems to be completely teachable with the set of simple lessons we currently have implemented. It is globally

complete to the extent that the viewpoints represented include all of the viewpoints routinely held by experts and any others that are important to learning the domain content. Finally, it is relatively *consistent*, in that the prerequisite relationships all run in the same direction. There are no cases where Lesson X is prerequisite to Lesson Y from one point of view while Lesson Y is prerequisite to Lesson X from another.




Figure 6
Goal Lattice Layer for Resistor Networks Course

It is not inevitable that courses will have these three properties. For example, consider the sort of introductory psychology course suggested by most current textbooks. There is little sense of completeness; lessons are included simple because of marketing needs and current fashion. There is no sense in which the simple lessons at the bottom of the hierarchy represent the simple foundations of concepts. Similarly, there is little coherence; the differing viewpoints, social, behavioral, cognitive, clinical, do not cover the same set of underlying basic concepts. Finally, there may not even be consistency. From one viewpoint, it may be best to teach sensory physiology before teaching about mental imagery; from another point of view a reverse ordering may seem obvious.

It seems appropriate to advance, as a hypothesis for future research, that knowledge-driven instructional systems will work best and be most implementable for those courses which have more or less coherent, consistent, and complete goal structures.

3.3 The Metaissue Layer

The third layer of the proposed architecture for instructional knowledge is the metaissue layer. Once again, it is useful to recount some of our reasoning in deciding on the need for such a layer and on what should be in it. Initially, we were motivated by a single issue: the conflict between the sorts of data used by very good teachers to decide on how to proceed with a given student and the data and reasoning used by the few expert instructional systems that have been built. In giving students assignments, teachers tend to rely on very broad representation of aptitude combined with a detailed knowledge of the curriculum. A child is "a good student," "a fast learner," "good in math," and/or "on page 93 in the book." In contrast, intelligent instructional systems, as envisioned by artificial intelligence researchers (e.g., Burton & Brown, 1982; Goldstein & Carr, 1977) construct a detailed student model which represents the best guess about

exactly what a student does and does not know of the specific material targeted by the course, e.g., "borrowing across zero," "adding single-digit integers," etc.

There is at least modest evidence (Burton & Brown, 1982) that teachers cannot determine microscopic representations of student knowledge status nearly as well as intelligent computer systems can, so the first hope for an approach that would be sensitive to the detailed specifics of changing student knowledge is probably an intelligent instructional computer system. This prompted us to think about how to represent aptitude data. We were, in this thinking, heavily influenced by the object-oriented approach we were taking (this approach will be discussed in Section 3.4 below). Basically, we were led to the following point of view.




Figure 7
The Metaissue Layer

Attending to a specific aptitude or some other metaissue in shaping the activities presented for the trainee in any lesson is simply a special case of shaping a lesson according to a specific viewpoint. That is, just as we might attend to differences between series and parallel circuits in our resistor network tutor and expect some students to have trouble with parallel circuits even after they have mastered series circuits, we could attend to differences among students in, say, reading ability or verbal facility of the student, and thus tailor our teaching to each student's capabilities seen from the verbal-facility point of view. This has led us to the architecture shown in Figure 7, in which the metaissue layer is simply the collection of goal nodes that are the origins of various viewpoint hierarchies embedded within the curriculum lattice.

3.4 Lesson Objects

So far, we have rather mysteriously presented descriptions of various structures of knowledge, implying that it is organized into lessons that can be considered from a variety of different viewpoints. Also, we have suggested that, somehow, the lessons or subgoals of the curriculum are connected with a representation of the knowledge they are trying to teach. Nowhere have we said just what a lesson is, just what the structure within one of these graphs might be like. To this I now turn.

Our fundamental approach to designing architectures for intelligent instructional systems is object oriented. That is, we see the design task as one of specifying a set of intelligent fragments of computer program and then orchestrating the interactions among these fragments. This approach originated with Smalltalk (Goldberg & Robson, 1983), a language developed in the course of trying to determine ways in which powerful personal workstations could change education. In conventional computer programs, the primary means of controlling the order in which computations take place, the task discipline, is by the sequencing of instructions. In object-oriented programming, control is passed when objects, "entities that combine the properties of procedures and data

since they perform computations and save local state" (Stefik & Bobrow, 1986), send messages to other objects.

So, for example, the way in which our resistor network tutor might be started is for the student to point to a box on the screen that says *Start*. That box would actually be a menu operated by an object. The object would perhaps respond to the student's action by telling one of the metaissue nodes in the tutor to teach the student. For example, if the approach favored by the designer were to teach primarily the relevant laws of electrical circuits, the menu object might send the *Laws* object (see Figure 3) a message to teach everything for which it is responsible. The *Laws* object would, in turn, ask its first prerequisite object, *Kirchhoff'sLaw*, to act, and that object might in turn tell the *I + Parallel* object, which teaches that current sums across the branches of a parallel circuit, to act. The *I + Parallel* object would then send messages in turn to its two prerequisites, dealing with current summing over branches and with the notion that the current in a branch of a parallel network is always less than the current passing through the network as a whole, to teach their stuff. At each level, when one subgoal of a goal was satisfied, the next would be sent a message to act, and so on.

This requires that each object contain all the data and all of the methods needed to completely achieve the goal to which it corresponds. This is not as cumbersome as it may sound; it is not necessary for each object to be a complete instructional computer program. Rather, objects can "inherit" some of their methods from higher-level objects. So, for example, if there are many objects that should teach their content via an exploratory electrical circuits simulation environment, they can all have a pointer to a single higher-level object that includes the program for such a simulation. Each object using the simulation might specialize it either by setting the values of variables to which the simulation program refers or by including specializing information in a message it sends to the simulator when it invokes that approach. The object-oriented approach is

valuable largely because it provides for a clear, understandable, and flexible means of assuring that the goals an instructional program has to achieve are clearly delineated and clearly "tasked" by relevant pieces of program code, even if the content to be taught under various circumstances must be determined dynamically.

Table 2 lists the contents that each goal lattice object must have in the kinds of instructional systems we are currently building. The list is split into two parts.

Declarative knowledge is the data that an object must have (or be prepared to accumulate during interactions with the student), what it is able to know. **Procedural knowledge** is the set of methods or programs that the object must have, what it is able to do. Each entry represents a specific kind of knowledge that must be present either explicitly, by being included in the object, or implicitly, via a pointer to the knowledge as part of a "parent" object.

Table 2
The Contents of a Lesson Object

Declarative Knowledge

Variables that identify how a given lessons' goals relate to the goals of other lessons (i.e., which lessons are prerequisite to the current one).

Variables that identify how the knowledge a lesson is trying to teach relates to the knowledge other lessons are trying to teach (pointers to the knowledge layer).

Variables that represent the student's mastery of the knowledge the lesson intends to teach (the student model).

Procedural Knowledge

Functions (methods) that generate instructional interventions based upon the student model held by the given object, including both manipulations of the interactive learning environment (perhaps a simulated laboratory or a problem generator) and various forms of coaching or advising.

Functions that decide if the given object is to blame for problems that arise while other lesson objects for which it is prerequisite are in control.

The declarative knowledge must include knowledge that places the object in the curriculum lattice, showing what its prerequisites are and also which objects assume it as a prerequisite. It must also include a specification of the specific parts of the knowledge layer (the representation of the target knowledge to be taught by the course) it is responsible for teaching. From this representation, it can provide explanations to the student. Finally, there must be a student model, a representation of which pieces of the target knowledge for the lesson the student appears to know and of the certainty of that diagnostic information.

The procedural knowledge each object must have is of two primary types. First, an object must (again, either explicitly or implicitly) be able to teach its target knowledge. Second, it must be able to decide whether a student's failure to perform adequately in learning a lesson for which it is prerequisite might be due to inadequate learning of its target knowledge. We call this *blame-taking* capability. The idea is that if things are going poorly in a lesson, the object teaching that lesson might ask each of its prerequisite objects to find out whether what it was supposed to teach is what the student is missing. With the curriculum structures discussed so far, that would seem to imply something like readministering pieces of the lesson posttest. However, there are more interesting possibilities to consider.

The goal object lattice structure, as discussed so far, bears striking resemblance to current practice. If one looks at a current elementary or high school textbook, one finds that each topic, and the exercises associated with it, is treated only once. One problem with curricula that actually follow such books (and most teachers do) is that the conditions of applicability for pieces of fact and process that students are taught are never reliably delineated. Students too often form rules for carrying out problem solving that are perfect for passing a unit test but maladaptive in the long run. For example, a student in elementary school given a set of arithmetic word problems might learn that it

isn't really necessary to read or understand the problem. If one finds words like *altogether*, one adds the numbers stated in the problem; if one finds words like *less*, one subtracts the smaller number from the larger; etc. Scanlon and O'Shea (in press) found similarly superficial strategies for use of specific equations that had recently been taught.

The way to avoid, or at least eventually remediate, such superficial learning is to combine different types of problems, to assure that problems occur in a variety of superficial contexts, so that the successful cues for various actions are cues based upon deep understanding of the problem. Jeffrey Bonar has been working on an extension of the architecture I describe here that tries to do this. In essence, higher-level curriculum objects keep a list of lower level objects that have recently been taught. Occasionally, a few problems are created that require various unpredictable combinations of this knowledge for their solution. This forces the student to look more deeply at the problem situation and work, like an expert, from deep understanding rather than from surface appearances.

However, the blame-taking problem becomes somewhat different in such cases. Rather than giving individual lesson posttests, testing for prerequisite knowledge in the original limited context used to present it, one wants to determine which pieces of knowledge that appear to have been mastered do not generalize to new situations. This appears to require a strategy of responding to a student's failure to handle a complex, multi-lesson problem either by giving hints or by giving a simpler problem, so that some of the candidate knowledge generalization failures can be ruled out. Work on how to do this is still proceeding; I mention it only because it may help give a sense of the character of blame-taking processes that we envision.

4.0 Context-Specific Prerequisite Content

I turn now to the final issue I wish to address, the complex nature of the prerequisite relationship. In particular, I want to show that the domain knowledge for which a lesson object is responsible is specific to the curricular context in which the lesson is invoked. That is, the knowledge that should be presented by a lesson depends upon the context in which that lesson is taught. To do this, I must compare the core content of a lesson being taught for the first time with its remedial content.

4.1 The Core Content of a Lesson

The target knowledge of a lesson object can be thought of as a set of pointers from that object to nodes in the knowledge layer. The subset of the domain (or expert) knowledge layer defined by those nodes and their relations to other nodes will generally not have a sharp boundary, because expert knowledge is highly interconnected. As can be seen in Figure 8, some of the nodes in a lesson's target knowledge will be connected to each other, while others will be outlying orphans, whose operational meaning, that is, the set of connections from a concept to node to other concept nodes, is defined primarily outside of the target subset. I use the term core content to refer to the subset of a lesson's target knowledge that is coherent, in the sense that its nodes are interconnected, with relatively few connections from a node to others outside that subset.

When a lesson is taught initially, its core content should be presented. That is, a coherent subset of the knowledge subsumed under the lesson should be taught. The density of detail in that coherent subset can vary with the aptitudes of the learner. Some learners should be taught all of the core content explicitly, while others can be expected to make at least the most direct and obvious inferences. In either case, it is impossible to teach explicitly the fringes of the target knowledge without also introducing knowledge outside the target subset, so it can be assumed that these fringe

pieces of knowledge have not been taught during the lesson's initial presentation. If a course, whether presented by human or by machine, is well taught, then the fringe knowledge for one lesson will be covered in another overlapping lesson or will be so immediately inferable from what is taught that it is optimal to assume that the student will learn it. Part of the artistry of curriculum design is to split the knowledge to be taught into pieces that cover the total set of target knowledge with no more overlap than the usual student will require.




Figure 8
A Curriculum Object Subsumes a Region of the Knowledge Layer Including More than Core Content

Further, there must be some mechanism for verifying how well the student is doing at learning what he is taught explicitly and completely and at inferring what it was hoped he would infer. In the design my colleagues and I have been evolving, blame taking provides this mechanism. When a lesson is proceeding poorly, an effort is made to determine which prerequisite material was not well presented, and that material is then retaught. Further, arbitrary objects can be created dynamically from time to time whose task it is to compose problems that require synthesis of several different pieces of already-taught knowledge and to test the student with them. When such problems fail one or more of the lessons they were based on should be reconsidered (i.e., the instructor/machine should consider reteaching them).

4.2 The Content of a Remedial Lesson

When a lesson is retaught remedially, there will generally be information to guide the selection of content that should be emphasized. In contrast to the emphasis placed on core content when a lesson is originally taught, it is crucial to teach the knowledge that links the core content of the to-be-remediated lesson with the core content of the lesson whose failure produced the need for remediation. We can make this point clearer by resorting to a graphical representation.

Look at Figure 9, which represents the interface between the goal lattice layer and the domain knowledge layer. Its point is that prerequisites are usually only partially overlapped by their superordinate lessons' domain knowledge. For example, the projection in the knowledge layer of Lesson A (Region a), is only partly contained in the projection of Lesson C (Region c), for which it is prerequisite. It has a different overlap with the projection of Lesson B (Region b), for which it is also prerequisite.




Figure 9
Higher-Level Objects' Projections Do Not Necessarily Cover the Projections of their Children

We can now proceed to define what content of a lesson should be taught when it is remediated. Basically, the emphasis should be on the nodes in the overlap between the projection of the superordinate lesson that failed and the projection of the prerequisite lesson that has taken blame. Figure 10 illustrates the area of Goal A's content that should be taught when A has taken blame for the failure of superordinate Goal B. Further, this overlap region should not be trimmed completely to produce a coherent core. Rather, its connections into the prerequisite lesson should also be pursued during remediation, and perhaps also its connections into the superordinate lesson that failed. This gives us a clear distinction between the lesson as originally taught, which emphasized core content, and the lesson as remediated, which involves contextually relevant context.

I conclude by noting that this specification of what should be taught in remedial instruction is probably not foreign to the master teacher, who undoubtedly makes such determinations intuitively. However, instructional machines must have principles to




Figure 10
The Appropriate Content of a Lesson when First Taught Is Not the Same as When Remediated

guide their performance, and the principles just stated seem reasonable candidates for inclusion in the teaching knowledge of such machines. In implementing this approach, I am sure that other candidates will also emerge. Finally, I suspect that the concerns we have had will be worth bringing to the attention of new teachers, who may not have a good idea of what the differences are between good remediation and simple repetition of instruction that failed the first time.

I also note that I have not considered another principle that seems worthy, namely that instruction should build from strength. This is not because I disagree with that principle but rather because I have nothing new ready to say yet. A good instructional system, especially when remediating, will want to order the knowledge that is presented so that it builds from knowledge the student is known to have already. I hope that my colleagues and I will have something to say soon about how this should be done.

References

- Anderson, J. A., Boyle, C. F., Farrell, R., & Reiser, B. (1984). Cognitive principles in the design of tutors. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*. Boulder, CO: The Institute of Cognitive Science and the University of Colorado.
- Bonar, J. (1985). *Bite-sized intelligent tutoring*. Technical Report. Pittsburgh, PA: University of Pittsburgh, Learning Research and Development Center. University of Pittsburgh.
- Brown, J. S., & Burton, R. R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2, 155-192.
- Brown, J. S., & Van Lehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. *Cognitive Science*, 4, 379-426.
- Burton, R. R., & Brown, J. S. (1982). An investigation of computer coaching for informal learning activities. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 79-98). New York: Academic Press.
- Carey, S. (In press). Are children fundamentally different kinds of thinkers and learners than adults? In S. F. Chipman, J. W. Segal, & R Glaser (Eds.), *Thinking and learning skills: Vol. 2. Research and open questions*. Hillsdale, NJ: Erlbaum.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.
- Chi, M.T.H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
- Gagné, R. M. (1962). The acquisition of knowledge. *Psychological Review*, 69, 355-365.

Gagné, R. M. (1971). *Conditions of learning*. New York: Holt Rinehart & Winston.

Glaser, R. (1984). Education and thinking: The role of knowledge. *American Psychologist*, 39, 93-104.

Glaser, R., Lesgold, A., & Lajoie, S. (in press). Toward a Cognitive Theory for the Measurement of Achievement. In R. R. Ronning, J. Glover, J. C. Conley, & Witt, J. C. *The influence of cognitive psychology on testing*. Hillsdale, NJ: Erlbaum.

Goldberg, A., & Robson, D. (1983). *Smalltalk-80: The language and its implementation*. Reading, MA: Addison-Wesley.

Goldstein, I., & Carr, B. (1977, October). The computer as coach: An athletic paradigm for intellectual education. *Proceedings of the 1977 Annual Conference*, Association for Computing Machinery, Seattle, WA, 227-233.

de Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, 208, 1335-1342.

Lesgold, A. M., Lajoie, S., Eastman, R., Eggan, G., Gitomer, D., Glaser, R., Greenberg, L., Logan, D., Magone, M., Weiner, A., Wolf, R., & Yengo, L. (1986, April). *Cognitive task analysis to enhance technical skills training and assessment*. Technical Report. Pittsburgh, PA: University of Pittsburgh, Learning Research and Development Center.

Scanlon, E., & O'Shea, T. (In press). Cognitive economy in physics reasoning: Implications for designing instructional materials. In H. Mandl & A. M. Lesgold (Eds.), *Learning issues for intelligent tutoring systems*. NY: Springer.

Shaughnessy, M. (1977). *Errors and expectations*. New York: Oxford.

Stefik, M., & Bobrow, D. (1986). Object-oriented programming: Themes and variations.
AI Magazine, 6, 40-62.

Van Lehn, K. (1983). On the representation of procedures in repair theory. In H. P. Ginsburg (Ed.), *The development of mathematical thinking*. (pp. 197-252). New York: Academic Press.

Young, R. M. & O'Shea, T. (1981). Errors in children's subtraction. *Cognitive Science*, 5, 152-177.

Footnotes

1I thank David Merrill for making this clear to me in his comments after a presentation I made at an AERA meeting in 1983.

2I mean to address both school instruction and technical training needs with the ideas presented in this essay. For ease of exposition, I shall use the term training to refer to both of these activities.

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Phillip L. Ackerman
University of Minnesota
Department of Psychology
Minneapolis, MN 55455

Dr. Beth Adelson
Department of Computer Science
Tufts University
Medford, MA 02155

Air Force Human Resources Lab
AFHRL/MPD
Brooks AFB, TX 78235

AFOSR,
Life Sciences Directorate
Bolling Air Force Base
Washington, DC 20332

Technical Director,
Army Human Engineering Lab
Aberdeen Proving Ground
MD 21005

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
Naval Training Systems Center
Orlando, FL 32813

Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. John Allen
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. James Anderson
Brown University
Center for Neural Science
Providence, RI 02912

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Nancy S. Anderson
Department of Psychology
University of Maryland
College Park, MD 20742

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Alan Baddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EE
ENGLAND

Dr. Patricia Baggett
University of Colorado
Department of Psychology
Box 345
Boulder, CO 80309

Dr. Eva L. Baker
UCLA Center for the Study
of Evaluation
145 Moore Hall
University of California
Los Angeles, CA 90024

Dr. Meryl S. Baker
Navy Personnel R&D Center
San Diego, CA 92152-6800

prof. dott. Bruno G. Bara
Unita di ricerca di
intelligenza artificiale
Universita di Milano
20122 Milano - via F. Sforza 23
ITALY

Dr. William M. Bart
University of Minnesota
Dept. of Educ. Psychology
330 Burton Hall
178 Pillsbury Dr., S.E.
Minneapolis, MN 55455

Dr. Jackson Beatty
Department of Psychology
University of California
Los Angeles, CA 90024

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. John Black
Teachers College
Columbia University
525 West 121st Street
New York, NY 10027

Dr. Jeff Bonar
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Robert Breaux
Code N-095R
Naval Training Systems Center
Orlando, FL 32813

Commanding Officer
CAPT Lorin W. Brown
NROTC Unit
Illinois Institute of Technology
3300 S. Federal Street
Chicago, IL 60616-3793

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Maj. Hugh Burns
AFHRL/IDE
Lowry AFB, CO 80230-5000

Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

LCDR Robert Carter
Office of the Chief
of Naval Operations
OP-01B
Pentagon
Washington, DC 20350-2000

Chair, Department of
Psychology
College of Arts and Sciences
Catholic University of
America
Washington, DC 20064

Dr. Michelene Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. L. J. Chmura
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Mr. Raymond E. Christal
AFHRL/MOE
Brooks AFB, TX 78235

Assistant Chief of Staff
for Research, Development,
Test, and Evaluation
Naval Education and
Training Command (N-5)
NAS Pensacola, FL 32508

Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Stanley Collyer
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Brian Dallman
3400 TTW/TTGXS
Lowry AFB, CO 80230-5000

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Diane Damos
Arizona State University
Department of Psychology
Tempe, AZ 85287

Dr. Denise Dellarosa
Department of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Dr. Stephanie Doan
Code 6021
Naval Air Development Center
Warminster, PA 18974-5000

Dr. Emanuel Donchin
University of Illinois
Department of Psychology
Champaign, IL 61820

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Susan Embretson
University of Kansas
Psychology Department
426 Fraser
Lawrence, KS 66045

Dr. Randy Engle
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. William Epstein
University of Wisconsin
W. J. Brodgen Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Martha Farah
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Beatrice J. Farr
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Marshall J. Farr
Farr-Sight Co.
2520 North Vernon Street
Arlington, VA 22207

Dr. Paul Feltovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3926
Springfield, IL 62708

Dr. Craig I. Fields
ARPA
1400 Wilson Blvd.
Arlington, VA 22209

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Alfred R. Fregly
AFOSR/NL
Bolling AFB, DC 20332

Dr. Michael Friendly
Psychology Department
York University
Toronto ONT
CANADA M3J 1P3

Julie A. Gadsden
Information Technology
Applications Division
Admiralty Research Establishment
Portsmouth, Portsmouth PO6 4AA
UNITED KINGDOM

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Dr. Lee Giles
AFOSR
Bolling AFB
Washington, DC 20332

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Marvin D. Glock
13 Stone Hall
Cornell University
Ithaca, NY 14853

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Daniel Gopher
Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL

Dr. Sherrie Gott
AFHRL/MODJ
Brooks AFB, TX 78235

Dr. T. Govindaraj
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332

Dr. Richard H. Granger
Department of Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. James G. Greeno
University of California
Berkeley, CA 94720

Dr. Henry M. Halff
Halff Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Bruce Hamill
The Johns Hopkins University
Applied Physics Laboratory
Laurel, MD 20707

Dr. John M. Hammer
Center for Man-Machine
Systems Research
Georgia Institute of Technology
Atlanta, GA 30332

Dr. Ray Hannapel
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Dr. Harold Hawkins
Office of Naval Research
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Dr. Frederick Hayes-Roth
Teknowledge
525 University Ave.
Palo Alto, CA 94301

Dr. Joan I. Heller
505 Haddon Road
Oakland, CA 94606

Dr. Geoffrey Hinton
Carnegie-Mellon University
Computer Science Department
Pittsburgh, PA 15213

Dr. James D. Hollan
MCC,
Human Interface Program
3500 West Balcones Center Dr.
Austin, TX 78759

Dr. John Holland
University of Michigan
2313 East Engineering
Ann Arbor, MI 48109

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Robert W. Holt
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Ms. Julia S. Hough
Lawrence Erlbaum Associates
8012 Greene Street
Philadelphia, PA 19144

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Earl Hunt
Department of Psychology
University of Washington
Seattle, WA 98105

Dr. Ed Hutchins
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Janet Jackson
Rijksuniversiteit Groningen
Biologisch Centrum, Vleugel D
Kerklaan 30, 9751 NN Haren (Gn.)
NETHERLANDS

Dr. R. J. K. Jacob
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Dr. Zachary Jacobson
Bureau of Management Consulting
365 Laurier Avenue West
Ottawa, Ontario K1A 0S5
CANADA

Pharm.-Chim. en Chef Jean Jacq
Division de Psychologie
Centre de Recherches du
Service de Sante des Armees
108 Boulevard Pinel
69272 Lyon Cedex 03, FRANCE

Dr. Robert Jannarone
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Claude Janvier
Directeur, CIRADE
Universite' du Quebec a Montreal
P.O. Box 8888, St. "A"
Montreal, Quebec H3C 3P8
CANADA

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

COL Dennis W. Jarvi
Commander
AFHRL
Brooks AFB, TX 78235-5601

Dr. Robin Jeffries
Hewlett-Packard Laboratories
P.O. Box 10490
Palo Alto, CA 94303-0971

Dr. Douglas H. Jones
Thatcher Jones Associates
P.O. Box 6640
10 Trafalgar Court
Lawrenceville, NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
Department of Psychology
University of California
Berkeley, CA 94720

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr.
P.O. Box 218
Yorktown Heights, NY 10598

Dr. David Kieras
University of Michigan
Technical Communication
College of Engineering
1223 E. Engineering Building
Ann Arbor, MI 48109

Dr. Walter Kintsch
Department of Psychology
University of Colorado
Campus Box 345
Boulder, CO 80302

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Mr. Al Kleider
Army Research Office
P.O. Box 12211
Research Triangle Park
North Carolina 27709-2211

Dr. Ronald Knoll
Bell Laboratories
Murray Hill, NJ 07974

Dr. Stephen Kosslyn
Harvard University
1238 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Dr. Kenneth Kotovsky
Department of Psychology
Community College of
Allegheny County
800 Allegheny Avenue
Pittsburgh, PA 15233

Dr. David H. Krantz
2 Washington Square Village
Apt. # 15J
New York, NY 10012

Dr. Patrick Kyllonen
325 Aderhold
Department of Educational
Psychology
University of Georgia
Athens, GA 30602

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152-6800

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. R. W. Lawler
ARI 6 S 10
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Alan M. Lesgold
Learning Research and
Development Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Alan Leshner
Deputy Division Director
Behavioral and Neural Sciences
National Science Foundation
1800 G Street
Washington, DC 20550

Dr. Jim Levin
Department of
Educational Psychology
210 Education Building
1310 South Sixth Street
Champaign, IL 61820-6990

Dr. John Levine
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309

Matt Lewis
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Library,
Naval War College
Newport, RI 02940

Library,
Naval Training Systems
Center
Orlando, FL 32813

Science and Technology Division,
Library of Congress
Washington, DC 20540

Dr. Jane Malin
Mail Code SR 111
NASA Johnson Space Center
Houston, TX 77058

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Dr. Humberto Maturana
University of Chile
Santiago
CHILE

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. James McBride
Psychological Corporation
c/o Harcourt, Brace,
Javanovich Inc.
1250 West 6th Street
San Diego, CA 92101

Dr. James L. McGaugh
Center for the Neurobiology
of Learning and Memory
University of California, Irvine
Irvine, CA 92717

Dr. Gail McKoon
CAS/Psychology
Northwestern University
1859 Sheridan Road
Kresge #230
Evanston, IL 60201

Dr. Joe McLachlan
Navy Personnel R&D Center
San Diego, CA 92152-6800

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. James S. McMichael
Navy Personnel Research
and Development Center
Code 05
San Diego, CA 92152

Dr. Barbara Means
Human Resources
Research Organization
1100 South Washington
Alexandria, VA 22314

Dr. Douglas L. Medin
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

Dr. George A. Miller
Department of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Dr. Andrew R. Molnar
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Dr. Nancy Morris
Search Technology, Inc.
5550-A Peachtree Parkway
Technology Park/Summit
Norcross, GA 30092

Dr. Randy Mumaw
Program Manager
Training Research Division
HUMRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elena Ave., 4th Floor
Redondo Beach, CA 90277

Chair, Department of
Computer Science
U.S. Naval Academy
Annapolis, MD 21402

Chair, Department of
Systems Engineering
U.S. Naval Academy
Annapolis, MD 21402

Technical Director,
Navy Health Research Center
P.O. Box 85122
San Diego, CA 92138

Dr. Allen Newell
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Mary Jo Nissen
University of Minnesota
N218 Elliott Hall
Minneapolis, MN 55455

Dr. A. F. Norcio
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Dr. Donald A. Norman
Institute for Cognitive
Science C-015
University of California, San Diego
La Jolla, California 92093

Deputy Technical Director,
NPRDC Code 01A
San Diego, CA 92152-6800

Director, Training Laboratory,
NPRDC (Code 05)
San Diego, CA 92152-6800

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152-6800

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Director, Human Factors
& Organizational Systems Lab.
NPRDC (Code 07)
San Diego, CA 92152-6800

Fleet Support Office.
NPRDC (Code 301)
San Diego, CA 92152-6800

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Technical Director,
Navy Personnel R&D Center
San Diego, CA 92152-6800

Commanding Officer,
Naval Research Laboratory
Code 2627
Washington, DC 20390

Dr. Harold F. O'Neil, Jr.
School of Education - WPH 801
Department of Educational
Psychology & Technology
University of Southern California
Los Angeles, CA 90089-0031

Dr. Michael Oberlin
Naval Training Systems Center
Code 711
Orlando, FL 32813-7100

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Office of Naval Research,
Code 1142BI
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142
800 N. Quincy St.
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142PS
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research.
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000
(8 Copies)

Psychologist.
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

Special Assistant for Marine
Corps Matters.
ONR Code 00MC
800 N. Quincy St.
Arlington, VA 22217-5000

Psychologist.
Office of Naval Research
Liaison Office, Far East
APO San Francisco, CA 96503

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Douglas Pearse
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. James W. Pellegrino
University of California,
Santa Barbara
Department of Psychology
Santa Barbara, CA 93106

Dr. Virginia E. Pendergrass
Code 711
Naval Training Systems Center
Orlando, FL 32813-7100

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637

Military Assistant for Training and
Personnel Technology,
OUSD (R & E)
Room 3D129, The Pentagon
Washington, DC 20301-3080

Dr. Steven Pinker
Department of Psychology
E10-018
M.I.T.
Cambridge, MA 02139

Dr. Martha Polson
Department of Psychology
Campus Box 346
University of Colorado
Boulder, CO 80309

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Michael I. Posner
Department of Neurology
Washington University
Medical School
St. Louis, MO 63110

Dr. Mary C. Potter
Department of Psychology
MIT (E-10-032)
Cambridge, MA 02139

Dr. Paul S. Rau
Code U-32
Naval Surface Weapons Center
White Oak Laboratory
Silver Spring, MD 20903

Dr. Lynne Reder
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. Wesley Regian
AFHRL/MOD
Brooks AFB, TX 78235

Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

Dr. Gil Ricard
Mail Stop C04-14
Grumman Aerospace Corp.
Bethpage, NY 11714

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. Paul R. Rosenbaum
Educational Testing Service
Princeton, NJ 08541

Dr. William B. Rouse
Search Technology, Inc.
5550-A Peachtree Parkway
Technology Park/Summit
Norcross, GA 30092

Dr. David Rumelhart
Center for Human
Information Processing
Univ. of California
La Jolla, CA 92093

Dr. Walter Schneider
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Miriam Schustack
Code 51
Navy Personnel R & D Center
San Diego, CA 92152-6800

Dr. Marc Sebrechts
Department of Psychology
Wesleyan University
Middletown, CT 06475

Dr. Colleen M. Seifert
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSF
La Jolla, CA 92093

Dr. Ben Schneiderman
Dept. of Computer Science
University of Maryland
College Park, MD 20742

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

LTCOL Robert Simpson
Defense Advanced Research
Projects Administration
1400 Wilson Blvd.
Arlington, VA 22209

Dr. H. Wallace Sinaiko
Manpower Research
and Advisory Services
Smithsonian Institution
801 North Pitt Street
Alexandria, VA 22314

Dr. Richard E. Snow
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

Dr. James J. Staszewski
Research Associate
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Robert Sternberg
Department of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Dr. Kurt Steuck
AFHRL/MOD
Brooks AFB
San Antonio TX 78235

Dr. Paul J. Sticha
Senior Staff Scientist
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

Dr. Kikumi Tatsuoka
CERL
252 Engineering Research
Laboratory
Urbana, IL 61801

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1185 Coleman Avenue, Box 580
Santa Clara, CA 95052

1987/03/09

Distribution List [Pittsburgh/Lesgold] NR 4422539

Dr. Sharon Tkacz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Douglas Towne
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

Headquarters, U. S. Marine Corps
Code MPI-20
Washington, DC 20380

Dr. William Uttal
NOSC, Hawaii Lab
Box 997
Kailua, HI 96734

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Beth Warren
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Keith T. Wescourt
FMC Corporation
Central Engineering Labs
1185 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Barbara White
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02238

Dr. Christopher Wickens
Department of Psychology
University of Illinois
Champaign, IL 61820

Dr. Heather Wild
Naval Air Development
Center
Code 6021
Warminster, PA 18974-5000

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Martin F. Wiskoff
Navy Personnel R & D Center
San Diego, CA 92152-6800

Dr. Dan Wolz
AFHRL/MOE
Brooks AFB, TX 78235

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Joe Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550

END

9-87

DTIC